

Reference Ontologies — Application Ontologies: Either/Or or Both/And?

Christopher Menzel
Texas A&M University
cmenzel@tamu.edu

The distinction between reference ontologies and application ontologies crept rather unobtrusively into the recent literature on knowledge engineering. A lot of the discourse surrounding this distinction – notably, the one framing the workshop generating this collection of papers – suggests the two types of ontologies are in some sort of opposition to one another. Thus, Borge *et al.* [3] characterize reference ontologies (more recently, *foundational* ontologies) as rich, axiomatic theories whose focus is to clarify the intended meanings of terms used in specific domains. Application ontologies, by contrast, provide a minimal terminological structure to fit the needs of a specific community. Reflecting their minimal nature, Masolo *et al.* [7] refer to such ontologies as “lightweight” ontologies. An application ontology can be lightweight in a second respect as well, namely, that it may not necessarily take the form of fully-fledged axiomatic theory. Rather, it might only be a taxonomy of the relevant domain, a division of the domain into a salient collection of classes, perhaps ordered by the subclass relation. Importantly, though, for an application ontology to “fit the needs of a specific community” needn’t require representational accuracy. In the “worst” case (from a reference ontology perspective), to fit the needs of a community is just to represent uncritically what people in that community think about the ontology’s domain.

The preceding paragraph is perhaps a good first cut, but it strikes me that the distinction between reference ontologies and application ontologies has not been clearly made. To clarify the distinction and its significance is my goal in this short paper. I conclude that, in fact, the distinction is not really an opposition; rather, reference ontologies and application ontologies reflect different aspects of a single methodology for ontology development.

1 Reference Ontologies

There appear to be three central characteristics of reference ontologies (ROs). We examine these in turn.

Theoretical Focus on representation The first characteristic of ROs is their theoretical focus on representation. ROs are constructed without any particular concerns for computational efficiency. Consequently, ROs avail themselves of (at least) the language of full first-order logic. Specifically, ROs avail themselves of:

- Arbitrary n -place predicates;

- Full classical negation;
- Unbounded, arbitrarily nested quantifiers.

The focus of ROs on representation is most clearly indicated in their generally unapologetic use of full first-order languages. The three features above are particularly noteworthy, as unrestricted use of any of them can render complete deductive procedures intractable, even undecidable.

Philosophical inclination toward realism The second feature of ROs is that their inclination toward philosophical realism. There are generally two elements of this realism:

- Metaphysical realism;
- Epistemological realism.

According to metaphysical realism, the World (Reality, What There Is) exists objectively in itself, independent of any mind. According to epistemological realism, the World is knowable by us. Thus, the philosophical standpoint underlying most ROs is that the World and its properties are *there* to be discovered. This implies, in turn, that the World, being objective and knowable, puts constraints on what we can say about it. Thus, in our ontologies can get it *wrong*. An RO is right just insofar as it accurately reflects, as far as it goes, the way the World is. This leads to our third feature of ROs.

Methodological emphasis on Truth Because our ROs can be wrong, there is in the construction of an RO a good reason to place a strong methodological emphasis on Truth. This has two practical implications:

- The central *function* of an ontology is to represent the World accurately and comprehensively; hence:
- The *quality* of an ontology a function of its accuracy and comprehensiveness.

ROs are all about getting the world — or some important piece of it — *right*. An ontology of time purports to describe its actual nature, to proffer the sober metaphysical truth on such matters as whether time is discrete, continuous, some combination of the two; whether there are timepoints or intervals, or both, and so on. Consequently, the quality of an ontology is judged along two dimensions: its accuracy — i.e., whether what it purports to be the case is in fact the case — and its comprehensiveness — i.e., whether it takes in a sufficiently broad spectrum of facts as to be significant.

2 Application Ontologies

Corresponding to our three features of reference ontologies are three salient features of application ontologies (AOs).

Theoretical Focus on Reasoning Unlike ROs, AOs are typically designed with some sort of computational application — and hence its attendant expressive limitations — in mind. Consequently, AOs are usually expressed in the language of some computationally tractable sublogic of full first-order logic (see, e.g., [6]). Such languages typically support:

- Reasoning about classes and “slots” through the use of unary and (limited) binary predicates;
- Conjunction and disjunction, but not negation;
- Limited quantification (e.g., wide-scope universal quantifiers only)

Philosophical inclination toward pragmatism/ instrumentalism/constructivism Unlike the strong realism underlying ROs, for AOs, take a far more pragmatic view of the world, both metaphysically and epistemologically. Specifically, the metaphysical presumption underlying a typical AO is the falsity, or at least the irrelevance, of metaphysical realism. The objects and structures we encounter in the world — those parts of it that matter to ontology, anyway — are social constructs, products of the evolving interaction between conscious, intelligent human agents and, at best, a substrate of unknowable *Dinge an sich*.

The corresponding epistemological presumption is that, even if metaphysical realism is true and there is an ultimate metaphysical reality to the world, that underlying reality probably unknowable anyway. Hence, what we can be said to know is simply what works.

Methodological emphasis on fidelity Methodologically, the central emphasis of an AO must be on *fidelity*, i.e., to be a *faithful expression* of the concepts/intuitions of relevant domain experts or sources. All that matters to an AO is how relevant domain experts *conceptualize* a given domain. The question of any sort *correspondence* between that conception and an objective external world is idle philosophical speculation with no bearing on the quality of the ontology, which is determined entirely by the extent of its fidelity.

On the face of it, these two approaches to ontology are profoundly different. However, the starkest differences are philosophical; indeed, those differences are probably irreconcilable. However, important as those differences might be conceptually, at the end of the day what we are engaged in is knowledge *engineering*. And as engineers, I suggest the following tendentious (not to say controversial) thesis: *the only components of the two approaches that ultimately matter are the theoretical and methodological*. These, I will argue are compatible, indeed complementary.

3 Pitfalls

An observation that suggests the complementarity of the two approaches is that, adopted separately, each risks serious pitfalls for which the other serves as a corrective.

3.1 Pitfalls of the Reference Ontology Approach

Realist Rigidity The first pitfall of the RO approach is *realist rigidity*. That is, the strong realist bias of this approach can deaden sensitivity to the conventional nature of “socially constructed” domains. Truth be told, it is difficult to reject a fairly strong form realism with regard to our best

physical theories. Intuitively, physics has gotten a lot of things pretty close to right, and has done so since Newton. As long as we restrict our attention to objects that are neither too massive, too small, nor moving too fast, Newton’s laws describe their physical properties with extraordinary precision. And, though things get admittedly dicier conceptually, contemporary relativity and quantum mechanics handle the rest with remarkable accuracy. *Surely*, it seems, the physicists have gotten things largely right.

However, the bulk of domains with which ontologies have to do are laden with objects that, while objective in some sense, can only be considered a species of human artifact — money, universities, email, laws, and so on. And the natures of these entities, *what they are* in a given domain, is to a large extent malleable. Consequently, there is often no single “correct” ontology for these domains, and hence often no way to reconcile related ontologies of these domains that have different origins.

The excellent paper by Vandenberghe *et al.* [11] in this volume makes this point with great cogency with respect to legal ontologies from different countries:

An ontology-based approach should enable the investigator to give the system a query in his language (“does the director of this company have previous convictions?”) and get a reliable answer in English that is affirmative if the “Vorstand” of the “Gesellschaft” has a “Vorstrafe”... While shared ontological commitment facilitates understanding, we encounter here a serious difference between similar projects in the natural sciences or medicine, and the legal domain. In the sciences, the underlying philosophical realism assumes convincingly that there is the possibility to agree on the relevant classification. The one shared reality sees to this. In law, the parliaments of the individual member states can create divergent realities. While German and UK law will agree that an extended prison sentence handed out by a court is a “previous conviction” and “Vorstrafe” respectively, it cannot be taken for granted that UK law recognises the equivalent of a German legal term such as “Strafbefehl”. If such a “Strafbefehl” is recorded against the name of the company director, should the system return an affirmative or a negative answer to our initial query?

The answer, of course, is that there is no objectively “correct” answer, only one that is determined “pragmatically based on the interests of [some] group of users with clearly defined goals” [REF, p 30].

Full First-order Obfuscation The second pitfall of the RO approach is *full first-order obfuscation*. Full FOL is a double-edged sword for the construction of ontologies. On the one hand, its great expressive power enables ontology developers to capture the ontological structure of a given domain in great detail and with exceptional clarity. The classic Ontolingua bibliography ontology [8], for example, is still an excellent example of a clear, well-structured ontology. A basic bibliographic item, or “BIBLIO-THING”, is elegantly defined in terms of a text of some ilk, an agent, an organization, a timepoint (the item’s publication date) a document, and a reference:

```
(subclass-partition BIBLIO-THING
  (setof biblio-text agent organization
    timepoint document reference))
```

The predicate “subclass-partition” in turn is clearly defined in terms of the predicates of the underlying KIF ontology, e.g.,

```
(iff (SUBCLASS-PARTITION ?C ?class-partition)
      (and (instance-of ?C class)
            (instance-of ?class-partition class-partition)
            (forall ?subclass
                    (=> (member ?subclass ?classpartition)
                        (subclass-of ?subclass ?C))))))
```

where, of course, the predicates here are either KIF primitive or defined ultimately in terms of KIF primitives.

On the other hand, it is relatively easy to write axioms in first-order logic, there is a danger of “over-first-orderization”, i.e., of trying to axiomatize too much. This can manifest itself in two ways. The first is rather innocuous: one can simply write more axioms than are necessary for one’s purposes in a given domain. The second is more serious: one can attempt to axiomatize domains that are not really amenable to formalization. This can give an illusion of clarity and rigor, when in fact it may just introduce subtle obscurity. The Standard Upper Merged Ontology (SUMO) [REF], for example, is an ambitious attempt to provide an axiomatization of a large class of higher level concepts to serve as a common conceptual base for ontological engineering. Many of the concepts are critical to this task — *class*, *relation*, *number*, *process* — and, of these, most are reasonably axiomatized. The axioms for many other concepts, however, obfuscate rather than illuminate. A pervasive problem with SUMO is its tendency to glibly axiomatize concepts that are notoriously controversial. Moreover, in many cases, these concepts are known to present serious theoretical challenges to axiomatization. SUMO essentially ignores these difficulties, and hence their axiomatizations of these concepts provide only an *illusion* of clarity and rigor where in fact there is little or none. For example, semantic notions like *content* and *information*, and intentional notions like *interpretation*, *agency*, and *belief* are widely-known to be exceedingly difficult and problematic to characterize axiomatically. SUMO blithely brings all five together in a single, rather astonishing axiom:

```
(=> (and (instance ?INTERPRET Interpreting)
          (agent ?INTERPRET ?AGENT)
          (patient ?INTERPRET ?CONTENT)
          (instance ?CONTENT ContentBearingObject))
      (exists (?PROP)
            (holdsDuring (EndFn (WhenFn ?INTERPRET))
                          (believes ?AGENT
                                    (containsInformation ?CONTENT ?PROP))))))
```

The apparent rigor here might lead one to think this axiom clarifies how these notions interact. However, it is useless without robust supporting axioms that give substance to the constituent concepts. SUMO’s supporting axioms are woefully inadequate. The *believes* relation, for instance, is axiomatized as a relation between agents and *formulas*, or more generally, linguistic things, a characterization widely recognized and acknowledged as untenable,¹ certainly not without a *con-*

¹Two well-known problems immediately come to mind. First, the approach is not obviously able to do justice to the intuition that people who speak different languages cannot have the same beliefs; since ‘*Es regnet*’ and ‘It is raining’

siderably richer account than SUMO provides. Again, we find that the *ContainsInformation* relation holds between a *ContentBearingObject* and a *Proposition*:

```
(domain containsInformation 1 ContentBearingObject)
(domain containsInformation 2 Proposition)
```

A notable special case of this is the *ContentBearingObject* is a sentence:

```
(=> (instance ?SENT Sentence)
      (exists (?PROP)
        (and (instance ?PROP Proposition)
              (containsInformation ?SENT ?PROP))))
```

Yet there is in SUMO no useful theory of propositions at all, no axioms about their nature and structure. Most critically, perhaps, there are no comprehension axioms of any sort that relate propositions to the sentences that allegedly express them in a way that enables us to reason meaningfully about them. For instance, SUMO tells us that both sentences and propositions have truth values, yet one cannot even infer from SUMO that if a sentence is true then the proposition it has as its content is true. And the frightfully difficult concept of truth itself (see, e.g., [9]), while occurring in many axioms, is not itself axiomatized even to the extent that one can prove that a proposition cannot be simultaneously true and false!

3.2 Pitfalls for the Application Ontology Approach

The AO approach faces significant theoretical and practical dangers of its own, notably, what I refer to as *pragmatic disarray* and *blindspots and monsters*.

Pragmatic Disarray As just noted, an excessive commitment to realism can lead to a certain sort of ontological inflexibility. But there is an equal and opposite danger that can arise from a studious avoidance of realism, namely, an overemphasis on the conventional element in human conceptualization in the formation of ontologies.

The construction of an effective ontology has both an empirical and a critical side. The empirical side involves careful attention to and recording of domain experts' characterizations of a given domain. The critical side involves anterior analysis of those conceptualizations for consistency and adequacy, and adequacy involves, at least, "getting things right" in some reasonably robust sense. If one rejects realism in all its forms, then it is difficult to justify putting any constraints (beyond mere consistency) on the acceptability of the conceptualizations of domain experts — the customer is always right. This in turn can lead to tolerance of poorly constructed ontologies that ultimately hinder the ontological enterprise.

are different linguistic items, a German speaker who believes the former believes something different than the English speaker standing next to her under an awning in a downpour. More seriously, taking linguistic entities to be the objects of intensional attitudes does not yield a satisfactory account of so-called *de re* beliefs, i.e., beliefs about specific individuals. Suppose there is someone whom John believes to be a spy. The natural SUMO representation of this is '(exists ?x (believes John (spy ?x)))'. But clearly, the content of John's attitude is not something involving the variable '?x', but some specific individual. Now *perhaps* the intuition SUMO's creators were trying to capture (the actual document notwithstanding) is that the formula '(spy ?x)' here is serving as a term denoting a proposition. But then we face the problem taken up next.

Blindspots and Monsters As just seen, use of highly expressive languages is no guarantee against either sloppiness or lack of content. However, such languages at least give the ontological engineer the *power*, if not the ability, to write rich and precise axioms. The expressively weak languages commonly used in the development of application ontologies, however, can encourage inaccuracy and sloppiness. Thus blindspots and monsters: the engineer can miss important classes and connections. This in turn could lead to the introduction of fictional, or at least underaxiomatized and hence largely useless, classes. Interestingly, despite the expressive capacities of its full-first order framework, SUMO suffers from this problem as well. When it moves away from higher-level concepts for which it borrowed heavily from existing, focused axiomatizations,² one finds only a rather overwhelming and disjointed array of axioms superficially relating one SUMO term to another. Where rich expressive capacity isn't even available, there is an even greater tendency toward the creation of this sort of superficial ontology.

4 Rapprochement

Pitfalls are, after all, avoidable. And as already suggested, the pitfalls of each our two approaches is generally best avoided by looking to the strengths of the other. Let us expand on this briefly in this final section.

4.1 “Getting It Right”

The philosophical approaches of both views share in common some notion of on ontology “getting it right.” For the realist, this means accuracy in representing an objective, external world. For the pragmatist, this means fidelity to the concepts and intuitions of domain experts and other relevant sources. These are really complementary sides of the same coin. As noted, only the most radical social constructionist denies the reality of an external world; at the same time, and only the most extreme realist would deny the human origins of a vast array of social concepts. The social world is in many ways conceptually malleable and contingent, and could have been constructed in very different ways. However, it could *not* have been constructed in *any* way we please. While the world does not *determine* the bulk of our social constructions, it nonetheless *constrains* them in very general, structural ways: How we characterize our mentality or categorize psychological dysfunction is in large measure our doing, but we cannot but see ourselves as conscious beings; there is surprising variability in how we can categorize the human body and its parts and functions, but we cannot but see ourselves as (at least, partly) material; we can't see the blood in our veins as a smooth liquid rather than a gas or solid; we cannot but see this material thing as larger or smaller than that other material thing; or (for normal perceivers) to see the world as colored rather than shades of gray; we could not have “constructed” light to be faster than it is, or gravity to repel rather than attract. Indeed, it is this underlying, objective structure, together with our own general capacities for rationality, creativity, and organization, that jointly give rise to the relatively stable and objective character of our laws, conventions, institutions and other social constructions. The realism of the

²In the interest of full disclosure, I should note that I am largely responsible for the axiomatization of the SUMO “structural” ontology of classes and relations. I believe it is reasonably well axiomatized, and I believe I can say this without much conceit, as most of the axioms are drawn from well known sources. Among the other reasonably well axiomatized theories in the SUMO, I include the mereotopology ontology, which borrows liberally from [10], [1], and [2], and the ontology of holes, taken largely from [4].

RO approach therefore undergirds and supports the robust pragmatism of the AO approach. We can unabashedly embrace both.

4.2 Complementarity and Expressive Power

The representational frameworks of the two approaches complement one another nicely as well. We found opposite and equal dangers: for the RO approach, obfuscation and the illusion of clarity; for the AO approach, blindspots and monsters. Obfuscation arises when one adds unnecessary (and perhaps inaccurate) complexity and detail. The weaker frameworks of the AO approach force one toward a measured simplicity — say all that you need to say, but no more.

On the other hand, both blindspots and monsters occur when one misses or misrepresents information about a domain that should have been captured. The conscientious knowledge engineer is compelled to express content as richly as possible. But even a conscientious engineer will not typically be as clear about the classes and logical connections in a domain as she would be were she using a more expressive representational framework. With a full first-order language at one's disposal, such an engineer is forced to think through different representational options at a level of detail that is impossible in the (far) weaker, computationally tractable languages of the AO approach.

The two representational frameworks can thus be complementary; thinking in terms of the the expressively weaker systems helps to keep one using a stronger framework focused; similarly, use of the stronger frameworks pushes one using a weaker framework toward greater accuracy and comprehensiveness.

4.3 A Balanced Methodology

The reflections above yield a more balanced methodology than is found in the RO and AO approaches alone. The AO is right to focus on the concepts and intuitions of domain experts, but the *justification* for that focus requires the realism of the RO approach. Domain experts are experienced in their domains. Through that experience they have become attuned not only to the interplay of their own concepts with respect to that domain, but also to the objective constraints on those concepts imposed by the world. Hence, the concepts of the domain experts are likely to be appropriate to and effective in that domain, i.e., to be “true” in the pragmatists sense. At the same time, however, the RO approach also tempers the AO here: because there is an objective reality underlying the concepts of the domain experts, there is room for *critique* — the experts might be wrong, or at a minimum their concepts less effective vis-a-vis the underlying reality than they might be otherwise. To cite a simple example, categorizations of mental dysfunction have changed dramatically over the past century in ways that are quite obviously more effective for the treatment of mental illness. The simple and most straightforward explanation of this is simply that modern categorizations better reflect the underlying structure of human mentality.

On the representational side, every application ontology should begin as a reference ontology. That is, the first order of business must be to express the structure of the domain at hand without any expressive constraints to help ensure accuracy and completeness. At the same time, reference ontology development proceeds always with an eye toward the application ontology and the corresponding expressive limitations of its computationally tractable framework, thereby hindering the tendency toward full first-order obfuscation.

5 Conclusion

In conclusion then, our two approaches to ontology development are actually two sides of a single methodology. Any opposition between the two approaches arises from the error of considering either approach to be complete in itself. Considered together, their complementarity is evident. The realism of the RO approach both tempers, and is tempered by, the social constructionism of the AO approach: many of the concepts in a given ontology are mind-dependent, social constructions, but constructions emerging out of knowable, mind-independent, underlying reality. Similarly, we can only hope to reason automatically within the computationally tractable representational framework of an AO ontology. But we best ensure the effectiveness of an application ontology by drawing it out of a comprehensive, overarching full first-order reference ontology rather than constructing it from scratch within its own expressively limited framework.

References

- [1] Borgo, S., Guarino, N., and Masolo, C., “A Pointless Theory of Space Based on Strong Connection and Congruence,” in L. C. Aiello and J. Doyle (eds.), *Principles of Knowledge Representation and Reasoning*, Morgan Kaufmann, 1996.
- [2] Borgo, S., Guarino, N., and Masolo, C., “An Ontological Theory of Physical Objects,” in L. Ironi (ed.), *Proceedings of the Eleventh International Workshop on Qualitative Reasoning*, Cortona (Italy), 1997, 223-231.
- [3] Borgo, S., Gangemi, A., Guarino, N., Masolo, C., and Oltramari, A., “Ontology Roadmap,” URL=<http://wonderweb.semanticweb.org/deliverables/documents/D15.pdf>.
- [4] Casati, R., and Varzi, A., *Holes and Other Superficialities*, Cambridge, MIT Press, 1995.
- [5] Guarino, N., “Understanding, Building and Using Ontologies: A Commentary to ‘Using Explicit Ontologies in KBS Development’,” *International Journal of Human and Computer Studies*, **46** (1997), 293-310.
- [6] Horrocks, I., Patel-Schneider, P., and van Harmelen, F., “From SHIQ and RDF to OWL: The Making of a Web Ontology Language,” *Journal of Web Semantics* **1**(1), 2003. URL=<http://www.cs.vu.nl/~frank/postscript/JWS03.pdf>.
- [7] Masolo, C., Borgo, S., Gangemi, A., Guarino, N., Oltramari, A., Schneider, L., “The Wonder-Web Library of Foundational Ontologies,” Preliminary Report (ver. 2.0, 15-08-2002), Laboratory for Applied Ontology, URL=<http://www.loa-cnr.it/Papers/WonderWebD17V2.0.pdf>.
- [8] Gruber, T., “The Bibliography Ontology,” <http://ksl-web.stanford.edu/htw/dme/thermal-kb-tour/biblio-onto.lisp.html>.
- [9] Martin, R. L., *Recent Essays on Truth and the Liar Paradox*, Oxford: Clarendon Press, 1984.
- [10] Smith, B., “Mereotopology: A Theory of Parts and Boundaries,” *Data and Knowledge Engineering*, **20**, 287-303.

- [11] Vandenberghe, W., Schafer, B., and Kingston, J., “Ontology Modelling In The Legal Domain — Realism Without Revisionism?”, this volume.